The Earth System Grid Center for Enabling Technologies (ESG-CET):

Scaling the Earth System Grid to Petascale Data

Climate simulation data are now securely accessed, monitored, cataloged, transported, and distributed to the national and international climate community

Semi-Annual Progress Report for the Period
October 1, 2006 through March 31, 2007

Table of Contents

1 Executive Summary

This report, which summarizes work carried out by the ESG-CET during the period October 1, 2006 through March 31, 2007, includes discussion of overall progress, period goals, highlights, presentations, and collaborations. To learn more about our project, please visit the Earth System Grid website. In addition, this report will be forwarded to the DOE SciDAC project management, the Office of Biological and Environmental Research (OBER) project management, national and international stakeholders (e.g., the *Community Climate System Model (CCSM),* the *Intergovernmental Panel on Climate Change (IPCC) 5th Assessment Report (AR5),* the *Climate Science Computational End Station (CCES)*, etc.), and collaborators.

The ESG-CET executive committee consists of David Bernholdt, ORNL; Ian Foster, ANL; Don Middleton, NCAR; and Dean Williams, LLNL. The ESG-CET team is a collective of researchers and scientists with diverse domain knowledge, whose home institutions include seven laboratories (ANL, LANL, LBNL, LLNL, NCAR, ORNL, PMEL) and one university (ISI/USC); all work in close collaboration with the project's stakeholders and domain researchers and scientists.

1.1 Overall goal for this reporting period

During this semi-annual reporting period, the ESG-CET focused its efforts on significant challenges involving management, architecture and framework design, metadata, and defining components. Issues pertinent to realization of a global petascale distributed architecture and framework in today's Grid computing environment included the coordination and synchronized design of overall architectural plans, use-case scenarios for each of the scientific subprojects (e.g. IPCC AR4/AR5, CCES, and CCSM), and continuing efforts to support and enhance existing ESG services. In addition, where appropriate, collaborations were established with national and international groups in order to exchange key ideas and open-source technologies.

1.2 Highlights

Since the start of the new funding cycle, the ESG-CET collaboration has continued to support the operations of the ESG data portals at NCAR and LLNL. During this reporting period also, an ESG data portal to support the CCES effort was established at ORNL. Currently, the NCAR portal serves a community of 4100 registered users, and allows access to over 144TB of data. The LLNL portal recently registered its 1000th user and has recorded a prestigious scientific publications list reaching well over 300 cited journal articles. Since their first release, combined download totals of the two portals reached well over 200TB, corresponding to over 800,000 files.

1.2.1 LLNL ESG Portal Highlights

The LLNL ESG portal is providing the world's climate scientists with the most complete collection of simulation data from current-generation climate models. The data include both simulations of past climate and projections of the future climate in 12 experiments by 23 models from 13 countries. This ESG archive was initiated by World Climate Research Programme scientists working on coupled ocean-atmosphere climate modeling in response to a request from the IPCC to consolidate predictions made for its Fourth Assessment (AR4) report. In February 2007, a DOE press release also highlighted the US role in studying and addressing global climate change.

December 2006 was a transition period, as the AR4 data archived at LLNL was relocated to a larger storage facility capable of storing 1 petabyte on rotating disks, of which 200 TB have been allocated for ESG-CET use. The facility, known as the Green Data Oasis (GDO), is currently serving 33 TB of AR4 data to more than 1100 users. Since coming on line, the LLNL ESG portal has achieved an average daily download rate of 700 gigabytes a day.

1.2.2 NCAR ESG Portal Highlights

Last summer, the NCAR ESG portal was redeployed on a new server to allow direct access to TeraGrid resources, and became an official TeraGrid Science Gateway. This migration involved a partial re-architecturing of some of the data services accessible through the portal, which is split into two separate components: an ESG Gateway, deployed on the TeraGrid server, which provides the user interface; and an ESG Node, deployed within the NCAR security perimeter, which hosts the back-end data services. We expect that this modular architecture will be reused and upgraded while building a next-generation data portal and services for ESG-CET.

More recently, in response to user feedback, the ESG NCAR portal is supporting *scripted data download*, i.e. the capability to run Unix scripts (through the standard *wget* client, and by means of *digest* authentication) that allow easy, sequential download of multiple files from the ESG holdings directly to a user's desktop, thus avoiding the need to execute multiple clicks on a web browser. The new *wget* capability complements the already operational DataMover-Lite functionality, thus providing ESG users with two reliable choices for multi-file data transfer.

Finally, prototype work and research has been conducted in the areas of Single Sign On among ESG Gateways, in order to evaluate the feasibility of federation requirements for the future ESG-CET system of distributed data services. A demo testbed based on the Central Authentication Service middleware developed by JA-SIG [CAS] was set up and deployed, allowing common authentication for ESG-like web portals representing two federated ESG Gateways.

1.2.3 ORNL ESG Portal Highlights

ORNL, in collaboration with the LLNL team, has deployed an ESG portal to support the Leadership Computing Facility's Computational Climate End Station (CCES), currently focusing on the CCSM Climate-Land Model Intercomparison Project (C-LAMP). Installation of the portal was completed in December 2006, and in order to maximize future flexibility, it is now running on an independent server from the main ESG ORNL node. Initial data archival from the C-LAMP project has been delayed due to the unexpectedly long spin-up phase of the model run, as well as technical issues involved in running the latest development version of the CCSM with biogeochemistry and associated data models enabled. We are expecting to receive initial monthly mean data from 3 runs (Experiments 1.2, 1.3, and 1.4) by mid-April with larger quantities of diurnal-average data from the same experiments in subsequent weeks. After testing the portal with the new model output data (since the C-LAMP metadata extends what was previously used in ESG), we will begin publishing further C-LAMP data into the ESG, as well as registering C-LAMP project members.

1.2.4 LANL ESG Portal Highlights

The LANL portal underwent an operating system upgrade (i.e., to SUSE Linux 10.1) in order to better support portal and grid software. LANL continues to serve high-resolution ocean data via the ESG NCAR portal. On related follow-on work, LANL team members have been working with Paraview to do remote visualization from ORNL. They are running 16-processor visualization jobs at ORNL and sending the output to a LANL desktop. So far, the prototype seems to be working pretty well.

1.2.5 LBNL Storage Resource Manager Highlights

The Storage Resource Managers (SRMs) for implementing access to the mass storage systems and the portal disk cache, as well as the DataMover-Lite (DML) client used to download files to users, have been highly stable over the last six months. The only reported errors were related to time-outs due to slow downloads by the clients.

1.2.6 PMEL Product Delivery Services Highlights

The conversion of the Live Access Server (LAS) into a generalized workflow engine has advanced to a BETA level of completion during the past six months. Endpoint testing suites have been developed to exercise the engine's robustness, and it is now functioning with high reliability in quasi-operational test implementations. The LAS defines a service protocol for requesting "products" (typically visualizations, tables, and file subsets) and provides wrappers to encapsulate legacy and new applications as services that are linked together in a workflow. Full wrappers have been developed for the Ferret application, for relational database access, and for observations accessible through the OPeNDAP "DAPPER" conventions. A prototype wrapper has been developed for the CDAT application. Code and logic for access to "CF" grids (curvilinear in latitude and longitude dimensions, and hybrid in the vertical), which are essential to the ESG ocean modeling community, have been developed. This logic has been incorporated into LAS and into ALPHA-level server-side analysis capabilities that are accessible through extensions to the OPeNDAP protocol.

1.2.7 ANL Security, Data, and Services Highlights

ANL worked closely with the ESG Security team to analyze important use-case scenarios, define requirements, and investigate solutions for the ESG security environment. One important result was the identification of a Single Sign On (SSO) requirement for browser-based clients, of which a prototype was subsequently implemented in the ESG Portal. This SSO facility is an important component of the larger grid security architecture that will be defined and implemented over the coming year.

Data discussions also began that led to the start of an integration of GridFTP into OPeNDAP. Similarly, discussions concerning server-side processing have also started coalescing into a prototype services architecture that will work in conjunction with the LAS architecture described above in section 1.2.6.

2 Overall Progress

During this reporting period, progress was made in the key areas that are necessary to meet goals and objectives.

2.1 ESG-CET Management Plan

The ESG-CET project management plan was carefully designed to construct and build the next generation Earth System Grid infrastructure for petascale computing. The management plan entails:

- Developing a project plan, which includes defining goals and objectives;
- Specifying tasks and how they will be implemented through subgroups;
- Defining milestones and deliverables, driven mainly (but not entirely) by the IPCC AR5 strict timetable;

- Developing internal and external resource requirements, project communications, and collaborations; and
- Associating budgets and timelines for task completion.

The usually phases of a project management plan also are included: implementation, evaluation, support, and maintenance. To view the project management plan in its entirety, see the ESG-CET reports website.

2.2 Architecture Design

During the first months of the project, the ESG-CET team engaged in a detailed process of 1) re-visioning the architecture of the current deployed operational systems (at NCAR, LLNL and ORNL), 2) defining the architecture of the next-generation data portals and services based on expanded functionality requirements, and 3) planning an evolutionary path forward from the former to the latter. As the outcome of this process, a *federated ESG-CET architecture* was defined, which will provide interoperability among the current operational ESG systems, and will greatly enhance the overall functionality available to the users.

The future ESG-CET architecture will be based on three tiers of data services (visit the architecture section of the ESG-CET website to view the design):

- Tier 1: A set of Global ESG Services (partially exposed via a Global ESG Portal) will provide shared functionality across the overall ESG-CET federation. These services will include 1) user registration and management, 2) common security services for access control to the federated resources, 3) metadata services for describing and searching the massive data holdings, 4) a common set of notification services and registry, and 5) global monitoring services. In particular, all ESG-CET sites will share a common database, so that a user will only have to register once and maintain a single set of credentials in order to access resources across the whole system – although access to specific resources (e.g. IPCC data served by the LLNL site) will still have to be approved by the corresponding local ESG administrators. Also, the definition of global metadata services for search and discovery will guarantee a user's ability to find data of interest throughout the whole federation, independently of the specific site at which a search is begun.
- Tier 2: A limited number of ESG Data Gateways will be set up to act as data-request brokers for specific communities, possibly also providing specific enhanced functionality and/or documentation, and supplying much needed fault-tolerance to the overall system. Services deployed on a Gateway will include the user interface for searching and browsing metadata, for requesting data (including analysis and visualization) products, and for orchestrating complex workflows. We expect the software stack to be deployed at the Gateway level will require considerable expertise to install and maintain, and thus we envision these Gateways being operated directly by ESG-CET engineering staff. Initially, three ESG Gateways will be set up: one at LLNL to serve the IPCC AR5 needs, one at ORNL to serve the Computational Climate End Station project, and one at NCAR to serve the CCSM and PCM model communities (and possibly also others).
- Tier 3: The actual data holdings, and the back-end data services that will be used to access the data, will reside on a (potentially large) number of federated ESG Nodes. These will typically host a suite of data and metadata services that are necessary to publish the data onto the ESG system, and to execute data-product requests formulated through an ESG Gateway. Because researchers and engineers at local institutions with variable levels of expertise will operate ESG Nodes, the software stack to be deployed at an ESG Node will be kept to a minimum, and supplied with detailed and exhaustive documentation. A single ESG Gateway will serve data requests to many associated ESG

nodes: for example, as part of the next IPCC project, more than 20 institutions are expected to set up data nodes connected to the LLNL Gateway.

2.3 Metadata

The design of the metadata database is at the heart of the ESG system, since the metadata model underlies other major ESG components, especially the search and browse facilities and the publishing system. Consequently this is one of the first tasks to have been addressed. During this reporting period, we held several meetings to develop the detailed scope and structure of the schema design.

The current software focuses on gridded datasets generated from climate model experiments. The next version will retain that focus, while anticipating the need to expand the scope of data served to related subject areas, such as the assimilation of observations or the prediction of climate impacts via models. Similarly, ESG will continue to add value by supporting derived and virtual datasets.

We have begun the process of designing a new search capability based on the idea of "faceted classification" that will provide several important features. At a given point in the search, the end user will see search terms and categories that apply within the current context, and will be able to avoid queries that return empty result sets. Similarly, organizing metadata around facets will provide important flexibility, since the categories can be changed and updated without impacting the rest of the system.

We are working with related metadata projects to ensure consistency with emerging community standards. For example, we have invited members of the Earth System Curator (ESC) and Numerical Model Metadata (NMM) projects to participate in the design process, and we are exploring how the respective schemas intersect. Because both ESC and NMM place more emphasis on the point of view of the data-producer, they have developed schemas that allow a rich description of the structure of models and model components; in contrast, ESG takes the viewpoint of the end user of data. The union of these data models thus will provide a richer, more comprehensive database, and will ensure that ESG can interface with the software systems derived from ESC.

2.4 Federated Metadata

One of the challenges of realizing a global, petascale, distributed architecture for the next generation of ESG is the design of a federated system that allows our sites to publish data sets and their associated metadata. In the last six months, we have held several meetings in order to converge on an architectural design for a distributed, federated metadata system.

The common architecture will include a single master metadata catalog that is hosted at the Global services layer of ESG. In practice, this service may be deployed at a particular ESG gateway site (e.g., NCAR, ORNL, or LLNL), but logically the metadata master catalog is gateway-independent: all updates to metadata must be performed on this master catalog.

In addition, the contents of the master metadata catalog will be replicated periodically on each of the ESG gateway nodes. This will allow users to issue metadata queries at any gateway. Having multiple, replicated catalogs that can answer user queries can provide load balancing so that the master catalog does not become a performance bottleneck. Multiple catalogs also provide improved accessibility: if the master catalog is unavailable, then no metadata update operations can be performed. However, metadata queries can be satisfied at any gateway using the replicated metadata catalogs.

We also agreed on a typical scenario for metadata harvesting, by which metadata for newly generated ESG data sets are stored in a federated catalog. In our architecture, each ESG data node generates

metadata for the data sets that it publishes, and then it stores this metadata in a local catalog. This may be harvested by the master metadata catalog in two ways. Either the data node sends, or "pushes", its metadata to the master catalog, or the master catalog "pulls" metadata from local catalogs.

2.5 Cyber Security

Secure access to data and resources plays a crucial role in the ESG. The security model must secure data and resources, but without creating an undue burden for the users and administrators. As a result of two face-to-face discussions, and numerous AccessGrid meetings and email exchanges between the security team members, the security architecture was substantially redesigned and is starting to take shape. One of the new requirements that emerged was a browser-client Single Sign On (SSO) requirement of which a prototype implementation was created. This SSO will allow the ESG portal functionality to be split among multiple servers while enforcing only authenticated access by the browser-clients. Furthermore, it will allow ESG in the future to leverage the identity management systems of selected partner organizations such that users will only have to authenticate within their home-domain.

The applicability of a number of grid security technologies including OSG VOMS, Shibboleth, ACEGI, and GT-CAS are being investigated. Further use-case analysis should help us to choose the right technologies and mechanisms without adding unnecessary complexity on the deployment side. Starting with the reference point of the ESG portal interface, existing open source technology for enabling Single Sign On was shown to be relatively easy to use and integrate within the ACEGI infrastructure. Both of these were integrated into a demo ESG portal to provide an important prototype validation for the architecture. We plan to extend this architecture with credential repositories like MyProxy as well as centralized Security Services that provide a grid interface so that clients can operate easily and securely within the ESG infrastructure and circumvent the portal interface where so required. Other technologies such as VOMS and Shibboleth may still need to be considered in order to maintain compatibility with TeraGrid as the TG security infrastructure is upgraded during the coming year.

2.6 Product Services

The ESG-CET is intended to serve customers with a broad range of sophistication. These users will range from numerical modelers, who want access to "raw" model output files and verbatim subsets of model output, to climate impacts investigators, who want rapid access to these data without the complexities of model-specific coordinate systems, to those users who only want to quickly visualize the overall behaviors of models. The petascale nature of the ESG data holdings require that significant levels of data reduction take place at the server in order to satisfy these customers – both through straightforward subsetting and decimation and through specific analysis operations, such as the computing of spatio-temporal averages. In the ESG architecture, we refer to the steps that convert raw data into analysis results and visualizations as "product services".

Story boards and use-case scenarios have been developed and are still being actively debated in weekly ESG meetings in order to characterize and better understand optimal experiences for the different classes of users. The Live Access Server (LAS), for example, is being adapted into a generalized workflow engine for the production of ESG products as described in section 1.2.6. A service-oriented approach is being employed in order to make the architecture adaptable to the range of users and products that are required. A high-functioning system has been created; capable of generating a wide range of visualization types (1D and 2D visualizations with arbitrary service-controllable customizations, such as contour levels and palettes file subsets) and basic operations such as averaging and determining extrema.

The analysis capabilities have been incorporated also into ALPHA-level server-side functionality that is accessible through extensions to the OPeNDAP protocol.

2.7 DataMover-Lite

During this reporting period, the main two activities included further development of the DataMover-Lite (DML), and of the Storage Resource Manager (SRM) for ESG. Future plans for multiple file transport include:

1) Overcoming problems with time-outs when downloading files with DML through addition of a capability to release files through the portal. For each file successfully downloaded by the client, the DML issues a "release" call provided to the SRM that manages the portal's disk cache. There are two advantages to this approach: 1) there is no need for the portal to allocate space for all the requested files at once, allowing better support for multiple users; and 2) file-lifetimes can be set to be sufficiently large, yet avoiding clogging of the disk-cache.
2) Plans to instrument DML, so that its usage by clients (number of files, total volume) will be tracked. It will also track downloads that complete successfully vs. those that are aborted before completion.
3) Plans to use a dynamic disk cache policy in the SRM for files accessed through the portal. Current indications are that making some subset of the files permanent does not utilize the disk cache well; where over 50% of permanent files are unread. A dynamic policy will allow frequently accessed files to remain longer in the cache; making files dynamically permanent based on usage patterns.

3 Architectural Design Diagrams

The architectural design diagrams referenced in Section 2 of this document can be viewed on the ESG-CET website.

4 ESG-CET Group Meetings

The ESG-CET executive committee holds weekly conference calls each Tuesday at 11:00 am MDT. These meetings discuss priorities and issues that make up the agenda for the weekly project meetings held via the AccessGrid (AG) every Thursday at 1:00 pm MDT. At these meetings, the entire team discusses project goals, design and development issues, technology, timelines, and milestones. Given the need for more in-depth conversations and examination of work requirements, the following face-to-face meetings were held during this reporting period:

4.1 Kickoff ESG-CET All-hands Meeting

In September, the ESG-CET project held a two-day all-hands meeting at NCAR in Boulder, CO dedicated to building the collaboration subgroups, user requirements, and project management for the project.

4.2 ESG-CET Architecture and Federation Meeting

In October, the ESG-CET architecture and federation subgroup held a two-day meeting at NCAR in Boulder, CO to determine the architecture infrastructure required for the project.

4.3 ESG-CET Metadata Meeting

In January, the ESG-CET metadata subgroup held a two-day meeting at LBNL in Berkeley, CA dedicated to developing our metadata schema, federated metadata, and metadata collaborations with outside groups for the project.

4.4 ESG-CET Cyber-security Meeting

In February, the ESG-CET cyber-security subgroup held a meeting at NCAR in Boulder, CO to discuss our cyber-security infrastructure for the project.

5 Collaborations

To effectively build an infrastructure capable of dealing with petascale data management and analysis, we established connections with other funded DOE Office of Science SciDAC projects and programs at various meetings and workshops, such as the SciDAC PI Workshop held in Atlanta, GA. In particular, collaborations have been established with the following groups:

5.1 A Scalable and Extensible Earth System Model for Climate Change Science (SciDAC Climate Application)

Collaborating with the SciDAC science application project *A Scalable and Extensible Earth System Model for Climate Change Science* (SEESM), ESG-CET has agreed to address the challenges of enabling data management, discovery, access, and advanced data analysis for their community. In constant contact with the SEESM PI, John Drake, and co-PI, Phil Jones, ESG-CET will act as the liaison for this group with other SciDAC projects (such as VACET and IUSV).

5.2 Intergovernmental Panel on Climate Change Fifth Assessment Report (IPCC AR5, United Nations Environment Program)

Dean Williams serves as the liaison between the current AR4 Working Group I committee leaders, led by Jerry Meehl, and the future AR5 Working Group I committee, to be led by Ron Stouffer. In October, Dean met with Ron and others to discuss proposed scenarios/experiments for AR5 and to work out strict deadlines for the project.

5.3 Computational Climate End Station (ORNL Leadership Computing Facility)

David Bernholdt (ORNL) serves as liaison with the Computational Climate End Station at the Leadership Computing Facility at ORNL. He is in regular contact with the CCES Chief Scientist John Drake (ORNL) and the leader of the C-LAMP project, Forrest Hoffman (ORNL). As noted above, a portal has been established at ORNL for the C-LAMP project, and it is expected to being receiving data in April 2007.

5.4 Open Science Grid (OSG, SciDAC Physics SAP) and the Center for Enabling Distributed Petascale Science (CEDPS, SciDAC CS CET)

The three-way group of OSG, CEDPS, and ESG has been working together to define areas of interaction and complementary work. Don Middleton (NCAR) met with Thomas Ndousse and colleagues from OSG and CEDPS in a March half-day meeting at the San Diego Supercomputing Center (SDSC) in order to advance this activity.

5.5 Visualization and Analytics Center for Enabling Technologies (VACET, SciDAC Visualization CET)

LANL team member (Phil Jones) and LLNL team members (Bob Drach, and Dean Williams) participated in the VACET project-wide meeting held in Utah. Science projects supported by ESG and the SciDAC awarded climate application were discussed and plans for collaborations are in progress. Integration of the VACET product visTrails and an ESG supported client application is now under way. VisTrails utilizes a visualization pipeline architecture that provides infrastructure to combine existing client applications into one environment. A particularly powerful use of visTrails is parameter exploration.

5.6 Workshop (DOE Cybersecurity R&D Challenges for Open Science: Developing a Roadmap and Vision)

Representing ANL, ESG, and NCAR, and the science/data community in general, Middleton & Ananthakrishnan attended the DOE-sponsored Cybersecurity Workshop, which was held in January 2007 at AGU headquarters in Washington, D.C. Middleton provided a short position paper along with concerns and scenarios. Rachana presented the whitepaper "Trust-Root Provisioning and Validation Facilities" by Frank Siebenlist, Michael Helm, Rachana Ananthakrishnan, and Ian Foster, which discusses deployment concerns for DOE projects like ESG.

5.7 TeraGrid (NSF)

Don Middleton continued to participate in TeraGrid meetings and activities, including the extension of ESG to function as a science gateway utilizing TG resources. Middleton attended the January 2007 TG Data Workshop in San Diego, CA.

5.8 Community Access to Global Cloud Resolving Model Data And Analyses (SciDAC Climate SAP)

Don Middleton (NCAR) attended the NSF CMMAP (Center for Multi-scale Modeling of Atmospheric Processes) in Kauai, Hawaii, meeting with colleagues Karen Schuchardt (PNNL) and John Helley (SDSC). This group worked on plans to develop an overall program plan that was complimentary relative to our specific project goals across SciDAC-2 and NSF Center for Enabling Technology goals.

5.9 A Data Domain to Model Domain Conversion Package (DMCP) for Sparse Climate Related Process Measurements (SciDAC Climate SAP)

ANL and ORNL team members (Dan Fraser and David Bernholdt) met with members of the DMCP project in late February and early March to explore possibilities for collaboration. Since much of what was discussed would take us beyond the current scope and funding of ESG-CET, we are currently deciding how to proceed.

5.10 Scientific Data Management Center for Enabling Technology (SciDAC CS CET)

The experience with the DataMover-Lite (DML) client component in ESG, led us to consider its use in the SDM center for moving files from sites that have one-time-password (OTP) security or other highly secure systems. Rather than "pulling" files into the client's site as done in ESG, the SDM center will use a DML for "pushing" files out of highly secure sites to their destinations. The SDM center is

planning to use this technology for a combustion project as well as for part of the workflow in a fusion project.

5.11 North American Regional Climate Change Assessment Program (NARCCAP, multi-agency)

Don Middleton (NCAR) continued to serve in the co-PI role for this project which will host its data products via the ESG system. During this reporting period, we refined the project's data management plan and began to develop the required software and system infrastructure.

6 Outreach, Presentations and Posters

List of talks and posters presented during this time period:

6.1 Paper and Presentation: eScience'06

Ann Chervenak (ISI/USC) presented a peer-reviewed paper at eScience 2006 in Amsterdam, Netherlands. The complete citation is: Ann Chervenak, Jennifer M. Schopf, Laura Pearlman, Mei-Hui Su, Shishir Bharathi, Luca Cinquini, Mike D'Arcy, Neill Miller, and David Bernholdt, Monitoring the Earth System Grid with MDS4, in *Second IEEE International Conference on e-Science and Grid Computing (e-Science'06)*, page 69, Los Alamitos, CA, USA, 2006, IEEE Computer Society.

6.2 Presentation: VACET All-hands Meeting

Bob Drach (LLNL) gave an ESG-CET presentation at the VACET All-hands meeting held at the University of Utah in Salt Lake City, Utah.

6.3 Presentation: Boulder Data Group

Don Middleton (NCAR) presented a talk entitled, "The Earth System Grid: Delivering Petascale Climate Data to a Global Community" at the inaugural meeting of the new Boulder Data Group, held in Boulder, CO.

6.4 Presentation: SC2006

ESG-CET activities were featured at booths sponsored by ANL, Global Grid Forum, LBNL, NCAR, and ORNL in the Exhibit Hall of SC2006, Tampa Bay, FL.

6.5 WMO Information System (WIS)

Don Middleton (NCAR) serves on multiple WMO WIS committees, and provided updates and overviews on ESG-CET at an intercommission coordination group in Beijing, China, and again at the international "Extraordinary Assembly" of WMO, a technical conference for information technology, held in Seoul, Korea November 2007.

6.6 Presentation: DOE LLNL Program Review

Dean Williams (LLNL) presented ESG-CET at the DOE PCMDI Program Review held in November 2007 at LLNL, CA.

6.7 Presentation: Curator Workshop

Dean Williams (LLNL) presented ESG-CET information at the Curator Workshop on Metadata held at GFDL in Princeton, NJ.

6.8 Position Paper: DOE Cybersecurity Workshop

Representing ESG, NCAR, and the science/data community in general, Don Middleton (NCAR) presented a short position paper at the DOE-sponsored Cybersecurity Workshop, which was held in January 2007 at AGU headquarters in Washington, D.C.

6.9 Poster: SciDAC'07 PI Workshop

Dean Williams (LLNL) and Don Middleton (NCAR) presented a poster on ESG-CET at the SciDAC '07 Workshop held in Atlanta, GA.

6.10 Poster: 3rd WGNE Workshop on Systematic Errors in Climate and NWP Models

Dean Williams (LLNL) presented a poster on ESG-CET at the 3rd WGNE Workshop on Systematic Errors in Climate and NWP Models held in February 2007 in San Francisco, CA.